

PATENT APPLICATION
ATTORNEY DOCKET NO. SUN-P9324-MEG

5

COMMIT INSTRUCTION TO SUPPORT
TRANSACTIONAL PROGRAM EXECUTION

Inventors: Marc Tremblay, Quinn A. Jacobson and Shailender Chaudhry

15 **Related Applications**

This application hereby claims priority under 35 U.S.C. §119 to U.S. Provisional Patent Application No. 60/447,128, filed on 13 February 2003, entitled "Transactional Memory," by inventors Shailender Chaudhry, Marc Tremblay and Quinn Jacobson (Attorney Docket No. SUN-P9322PSP).

20 **[0001]** This application is also related to a non-provisional U.S. patent application entitled, "Start Transactional Execution (STE) Instruction to Support Transactional Program Execution," by inventors Shailender Chaudhry Marc Tremblay and Quinn A. Jacobson, filed on the same day as the instant application (Attorney Docket No. SUN-P9323-MEG).

25 **[0002]** This application is also related to a non-provisional U.S. patent application entitled, "Fail Instruction to Support Transactional Program Execution," by inventors Shailender Chaudhry Marc Tremblay and Quinn A. Jacobson, filed on the same day as the instant application (Attorney Docket No. SUN-P9325-MEG).

30

BACKGROUND

Field of the Invention

5 [0003] The present invention relates to techniques for improving the performance of computer systems. More specifically, the present invention relates to a method and an apparatus for avoiding the overhead involved in using locks by transactionally executing critical sections of code.

Related Art

10 [0004] Computer system designers are presently developing mechanisms to support multi-threading within the latest generation of Chip-Multiprocessors (CMPs) as well as more traditional Shared Memory Multiprocessors (SMPs). With proper hardware support, multi-threading can dramatically increase the performance of numerous applications. However, as microprocessor performance
15 continues to increase, the time spent synchronizing between threads (processes) is becoming a large fraction of overall execution time. In fact, as multi-threaded applications begin to use even more threads, this synchronization overhead becomes the dominant factor in limiting application performance.

20 [0005] From a programmer's perspective, synchronization is generally accomplished through the use locks. A lock is typically acquired before a thread enters a critical section of code, and is released after the thread exits the critical section. If another thread wants to enter the same critical section, it must acquire the same lock. If it is unable to acquire the lock, because a preceding thread has grabbed the lock, the thread must wait until the preceding thread releases the lock.
25 (Note that a lock can be implemented in a number of ways, such as through atomic operations or semaphores.)

5 [0006] Unfortunately, the process of acquiring a lock and the process of releasing a lock are very time-consuming in modern microprocessors. They involve atomic operations, which typically flush the load buffer and store buffer, and can consequently require hundreds, if not thousands, of processor cycles to complete.

10 [0007] Moreover, as multi-threaded applications use more threads, more locks are required. For example, if multiple threads need to access a shared data structure, it is impractical for performance reasons to use a single lock for the entire data structure. Instead, it is preferable to use multiple fine-grained locks to lock small portions of the data structure. This allows multiple threads to operate on different portions of the data structure in parallel. However, it also requires a single thread to acquire and release multiple locks in order to access different portions of the data structure.

15 [0008] In some cases, locks are used when they are not required. For example, many applications make use of “thread-safe” library routines that use locks to ensure that they are “thread-safe” for multi-threaded applications. Unfortunately, the overhead involved in acquiring and releasing these locks is still incurred, even when the thread-safe library routines are called by a single-threaded application.

20 [0009] Applications typically use locks to ensure mutual exclusion within critical sections of code. However, in many cases threads will not interfere with each other, even if they are allowed to execute a critical section simultaneously. In these cases, mutual exclusion is used to prevent the unlikely case in which threads actually interfere with each other. Consequently, in these cases, the overhead involved in acquiring and releasing locks is largely wasted.

25

[0010] Hence, what is needed is a method and an apparatus that reduces the overhead involved in manipulating locks when accessing critical sections of code.

5 [0011] Although compilers and special-purpose hardware can be developed to automatically solve some of the above-described problems, in many cases, it is desirable to provide instruction-level support to enable programmers and compilers to control such solutions.

[0012] Hence, what is needed is a method and an apparatus that provides instruction-level support to solve the above-described problems.

10

SUMMARY

[0013] One embodiment of the present invention provides a system that supports a start transactional execution (STE) instruction, wherein the STE instruction marks the beginning of a block of instructions to be executed transactionally. Upon encountering the STE instruction during execution of a
15 program, the system commences transactional execution of a block of instructions following the STE instruction. Changes made during this transactional execution are not committed to the architectural state of the processor until the transactional execution successfully completes (which means that it completes without
20 violating the memory consistency model or without other errors).

[0014] In a variation on this embodiment, the STE instruction specifies an action to take if transactional execution of the block of instructions fails.

[0015] In a further variation, the action to take can include branching to a location specified by the STE instruction.

25 [0016] In a further variation, the action to take can include acquiring a lock on the block of instructions.

[0017] In a further variation, the action to take can include setting state information within the processor to indicate a failure during transactional execution of the block of instructions. This enables other software executed by the processor to manage the failure.

5 [0018] In a variation on this embodiment, potentially interfering data accesses from other processes are allowed to proceed during the transactional execution of the block of instructions.

 [0019] In a variation on this embodiment, if the transactional execution completes without encountering an interfering data access from another process or
10 other type of failure, the system atomically commits changes made during the transactional execution to the architectural state of the processor, and resumes normal non-transactional execution.

 [0020] In a variation on this embodiment, if an interfering data access from another process is encountered during the transactional execution, the system
15 discards changes made during the transactional execution, and attempts to re-execute the block of instructions.

 [0021] In a variation on this embodiment, the block of instructions to be executed transactionally comprises a critical section.

 [0022] In a variation on this embodiment, commencing transactional
20 execution of the block of instructions involves saving the state of processor registers. It also involves configuring the processor to mark cache lines during loads and stores that take place during transactional execution. It additionally involves configuring the processor to continually monitor data references from other threads to detect interfering data references.

25 [0023] In a variation on this embodiment, the STE instruction is a native machine code instruction of the processor.

[0024] In a variation on this embodiment, the STE instruction is defined in a platform-independent programming language.

5 [0025] One embodiment of the present invention provides a system that facilitates executing a commit instruction, which marks the end of a block of instructions to be executed transactionally. Upon encountering the commit instruction during execution of a program, the system successfully completes transactional execution of the block of instructions preceding the commit instruction. Changes made during the transactional execution are not committed to the architectural state of the processor until the transactional execution
10 successfully completes.

[0026] In a variation on this embodiment, the system successfully completes the transactional execution by atomically committing changes made during the transactional execution, and resuming normal non-transactional execution.

15 [0027] In a further variation, while committing changes made during the transactional execution, the system commits register file changes made during transactional execution and clears load marks from cache lines. The system also treats store-marked cache lines as locked, thereby causing other processes to wait to access the store-marked cache lines. The system subsequently commits store
20 buffer entries generated during transactional execution to memory, which involves unmarking, and thereby unlocking, corresponding store-marked cache lines.

[0028] In a variation of the commit instruction, after committing changes made during the transactional execution, the system commences transactional execution of the block of instructions following the commit instruction.

25 [0029] In a variation on this embodiment, the commit instruction is a native machine code instruction of the processor.

[0030] In a variation on this embodiment, the commit instruction is defined in a platform-independent programming language.

[0031] One embodiment of the present invention provides a system that supports executing a fail instruction, which terminates transactional execution of a block of instructions. During operation, the system facilitates transactional execution of a block of instructions within a program, wherein changes made during the transactional execution are not committed to the architectural state of the processor until the transactional execution successfully completes. If a fail instruction is encountered during this transactional execution, the system terminates the transactional execution without committing results of the transactional execution to the architectural state of the processor.

[0032] In a variation on this embodiment, terminating the transactional execution involves discarding changes made during the transactional execution.

[0033] In a variation on this embodiment, while discarding changes made during the transactional execution, the system discards register file changes and clears load marks from cache lines. The system also drains store buffer entries generated during transactional execution, and clears store marks from cache lines.

[0034] In a variation on this embodiment, while terminating the transactional execution, the system branches to a location specified by a corresponding start transactional execution (STE) instruction.

[0035] In a variation on this embodiment, while terminating the transactional execution, the system branches to a location specified by the fail instruction.

[0036] In a variation on this embodiment, while terminating the transactional execution, the system attempts to re-execute the block of instructions.

[0037] In a variation on this embodiment, the fail instruction is a native machine code instruction of the processor.

[0038] In a variation on this embodiment, the fail instruction is defined in a platform-independent programming language.

5

BRIEF DESCRIPTION OF THE FIGURES

[0039] FIG. 1 illustrates a computer system in accordance with an embodiment of the present invention.

[0040] FIG. 2 illustrates how a critical section is executed in accordance with an embodiment of the present invention.

10

[0041] FIG. 3 presents a flow chart illustrating the transactional execution process in accordance with an embodiment of the present invention.

[0042] FIG. 4 presents a flow chart illustrating a start transactional execution (STE) operation in accordance with an embodiment of the present invention.

15

[0043] FIG. 5 presents a flow chart illustrating how load marking is performed during transactional execution in accordance with an embodiment of the present invention.

[0044] FIG. 6 presents a flow chart illustrating how store marking is performed during transactional execution in accordance with an embodiment of the present invention.

20

[0045] FIG. 7 presents a flow chart illustrating how a commit operation is performed in accordance with an embodiment of the present invention.

[0046] FIG. 8 presents a flow chart illustrating how changes are discarded after transactional execution completes unsuccessfully in accordance with an embodiment of the present invention.

25

[0047] FIG. 9 illustrates instructions to support transactional execution in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

5 [0048] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications
10 without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

15 Computer System

 [0049] FIG. 1 illustrates a computer system 100 in accordance with an embodiment of the present invention. Computer system 100 can generally include any type of computer system, including, but not limited to, a computer system based on a microprocessor, a mainframe computer, a digital signal processor, a
20 portable computing device, a personal organizer, a device controller, and a computational engine within an appliance. As is illustrated in FIG. 1, computer system 100 includes processors 101 and level 2 (L2) cache 120, which is coupled to main memory (not shown). Processor 102 is similar in structure to processor 101, so only processor 101 is described below.

25 [0050] Processor 101 has two register files 103 and 104, one of which is an “active register file” and the other of which is a backup “shadow register file.” In one embodiment of the present invention, processor 101 provides a flash copy

operation that instantly copies all of the values from register file 103 into register file 104. This facilitates a rapid register checkpointing operation to support transactional execution.

5 **[0051]** Processor 101 also includes one or more functional units, such as adder 107 and multiplier 108. These functional units are used in performing computational operations involving operands retrieved from register files 103 or 104. As in a conventional processor, load and store operations pass through load buffer 111 and store buffer 112.

10 **[0052]** Processor 101 additionally includes a level one (L1) data cache 115, which stores data items that are likely to be used by processor 101. Note that each line in L1 data cache 115 includes a “load marking bit,” which indicates that a data value from the line has been loaded during transactional execution. This load marking bit is used to determine whether any interfering memory references take place during transactional execution as is described below with reference to
15 FIGs. 3-8. Processor 101 also includes an L1 instruction cache (not shown).

[0053] Note that load marking does not necessarily have to take place in L1 data cache 115. In general load marking can take place at any level cache, such as L2 cache 120. However, for performance reasons, the load marking takes place at the cache level that is closest the processor as possible, which in this case
20 is L1 data cache 115. Otherwise, loads would have to go to L2 cache 120 even on an L1 hit.

[0054] L2 cache 120 operates in concert with L1 data cache 115 (and a corresponding L1 instruction cache) in processor 101, and with L1 data cache 117 (and a corresponding L1 instruction cache) in processor 102. Note that L2 cache
25 120 is associated with a coherency mechanism 122, such as the reverse directory structure described in U.S. Patent Application No. 10/186,118, entitled, “Method and Apparatus for Facilitating Speculative Loads in a Multiprocessor System,”

filed on June 26, 2002, by inventors Shailender Chaudhry and Marc Tremblay (Publication No. US-2002-0199066-A1). This coherency mechanism 122 maintains “copyback information” 121 for each cache line. This copyback information 121 facilitates sending a cache line from L2 cache 120 to a requesting
5 processor in cases where the current version of the cache line must first be retrieved from another processor.

[0055] Each line in L2 cache 120 includes a “store marking bit,” which indicates that a data value has been stored to the line during transactional execution. This store marking bit is used to determine whether any interfering
10 memory references take place during transactional execution as is described below with reference to FIGs. 3-8. Note that store marking does not necessarily have to take place in L2 cache 120.

[0056] Ideally, the store marking takes place in the cache level closest to the processor where cache lines are coherent. For write-through L1 data caches,
15 writes are automatically propagated to L2 cache 120. However, if an L1 data cache is a write-back cache, we perform store marking in the L1 data cache. (Note that the cache coherence protocol ensures that any other processor that subsequently modifies the same cache line will retrieve the cache line from the L1 cache, and will hence become aware of the store mark.)

20

Executing a Critical Section

[0057] FIG. 2 illustrates how a critical section is executed in accordance with an embodiment of the present invention. As is illustrated in the left-hand side of FIG. 2, a process that executes a critical section typically acquires a lock
25 associated with the critical section before entering the critical section. If the lock has been acquired by another process, the process may have to wait until the other process releases the lock. Upon leaving the critical section, the process releases

the lock. (Note that the terms “thread” and “process” are used interchangeably throughout this specification.)

[0058] A lock can be associated with a shared data structure. For example, before accessing a shared data structure, a process can acquire a lock on the shared data structure. The process can then execute a critical section of code that accesses the shared data structure. After the process is finished accessing the shared data structure, the process releases the lock.

[0059] In contrast, in the present invention, the process does not acquire a lock, but instead executes a start transactional execution (STE) instruction before entering the critical section. If the critical section is successfully completed without interference from other processes, the process performs a commit operation, to commit changes made during transactional execution. This sequence of events is described in more detail below with reference to FIGs. 3-8.

[0060] Note that in one embodiment of the present invention a compiler replaces lock-acquiring instructions with STE instructions, and also replaces corresponding lock releasing instructions with commit instructions. (Note that there may not be a one-to-one correspondence between replaced instructions. For example, a single lock acquisition operation comprised of multiple instructions may be replaced by a single STE instruction.) The above discussion presumes that the processor’s instruction set has been augmented to include an STE instruction and a commit instruction. These instructions are described in more detail below with reference to FIGs. 3-9.

Transactional Execution Process

[0061] FIG. 3 presents a flow chart illustrating how transactional execution takes place in accordance with an embodiment of the present invention. A process first executes an STE instruction prior to entering of a critical section

of code (step 302). Next, the system transactionally executes code within the critical section, without committing results of the transactional execution (step 304).

5 **[0062]** During this transactional execution, the system continually monitors data references made by other processes, and determines if an interfering data access (or other type of failure) takes place during transactional execution. If not, the system atomically commits all changes made during transactional execution (step 308) and then resumes normal non-transactional execution of the program past the critical section (step 310).

10 **[0063]** On the other hand, if an interfering data access is detected, the system discards changes made during the transactional execution (step 312), and attempts to re-execute the critical section (step 314).

15 **[0064]** In one embodiment of the present invention, the system attempts the transactionally re-execute the critical section zero, one, two or more times. If these attempts are not successful, the system reverts back to the conventional technique of acquiring a lock on the critical section before entering the critical section, and then releasing the lock after leaving the critical section.

20 **[0065]** Note that an interfering data access can include a store by another process to a cache line that has been load marked by the process. It can also include a load or a store by another process to a cache line that has been store marked by the process.

25 **[0066]** Also note that circuitry to detect interfering data accesses can be easily implemented by making minor modifications to conventional cache coherence circuitry. This conventional cache coherence circuitry presently generates signals indicating whether a given cache line has been accessed by another processor. Hence, these signals can be used to determine whether an interfering data access has taken place.

Starting Transactional Execution

[0067] FIG. 4 presents a flow chart illustrating a start transactional execution (STE) operation in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 302 of the flow chart in FIG. 3. The system starts by checkpointing the register file (step 402).
5 This can involve performing a flash copy operation from register file 103 to register file 104 (see FIG. 1). In addition to checkpointing register values, this flash copy can also checkpoint various state registers associated with the currently executing process. In general, the flash copy operation checkpoints enough state to
10 be able to restart the corresponding thread.

[0068] At the same time the register file is checkpointed, the STE operation also causes store buffer 112 to become “gated” (step 404). This allows existing entries in store buffer to propagate to the memory sub-system, but prevents new store buffer entries generated during transactional execution from
15 doing so.

[0069] The system then starts transactional execution (step 406), which involves load-marking and store-marking cache lines, if necessary, as well as monitoring data references in order to detect interfering references.

Load Marking Process

[0070] FIG. 5 presents a flow chart illustrating how load marking is performed during transactional execution in accordance with an embodiment of the present invention. During transactional execution of a critical section, the system performs a load operation. In performing this load operation if the load
25 operation has been identified as a load operation that needs to be load-marked, system first attempts to load a data item from L1 data cache 115 (step 502). If the load causes a cache hit, the system “load marks” the corresponding cache line in

L1 data cache 115 (step 506). This involves setting the load marking bit for the cache line. Otherwise, if the load causes a cache miss, the system retrieves the cache line from further levels of the memory hierarchy (step 508), and proceeds to step 506 to load mark the cache line in L1 data cache 115.

5

Store Marking Process

[0071] FIG. 6 presents a flow chart illustrating how store marking is performed during transactional execution in accordance with an embodiment of the present invention. During transactional execution of a critical section, the system performs a store operation. If this store operation has been identified as a store operation that needs to be store-marked, the system first prefetches a corresponding cache line for exclusive use (step 602). Note that this prefetch operation will do nothing if the line is already located in cache and is already in an exclusive use state.

10 [0072] Since in this example L1 data cache 115 is a write-through cache, the store operation propagates through L1 data cache 115 to L2 cache 120. The system then attempts to lock the cache line corresponding to the store operation in L2 data cache 115 (step 604). If the corresponding line is in L2 cache 120 (cache hit), the system “store marks” the corresponding cache line in L2 cache 120 (step 610). This involves setting the store marking bit for the cache line. Otherwise, if the corresponding line is not in L2 cache 120 (cache miss), the system retrieves the cache line from further levels of the memory hierarchy (step 608) and then proceeds to step 610 to store mark the cache line in L2 cache 120.

20 [0073] Next, after the cache line is store marked in step 610, the system enters the store data into an entry of the store buffer 112 (step 612). Note that this store data will remain in store buffer 112 until a subsequent commit operation

takes place, or until changes made during the transactional execution are discarded.

[0074] Note that a cache line that is store marked by a given thread can be read by other threads. Note that this may cause the given thread to fail while the other threads continue.

Commit Operation

[0075] FIG. 7 presents a flow chart illustrating how a commit operation is performed after transactional execution completes successfully in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 308 of the flow chart in FIG. 3.

[0076] The system starts by treating store-marked cache lines as though they are locked (step 702). This means other processes that request a store-marked line must wait until the line is no longer locked before they can access the line. This is similar to how lines are locked in conventional caches.

[0077] Next, the system clears load marks from L1 data cache (step 704).

[0078] The system then commits entries from store buffer for stores that are identified as needing to be marked, which were generated during the transactional execution, into the memory hierarchy (step 706). As each entry is committed, a corresponding line in L2 cache is unlocked.

[0079] The system also commits register file changes (step 708). For example, this can involve functionally performing a flash copy between register file 103 and register file 104 in the system illustrated in FIG. 1.

Discarding Changes

[0080] FIG. 8 presents a flow chart illustrating how changes are discarded after transactional execution completes unsuccessfully in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 312 of the flow chart in FIG. 3. The system first discards register file changes made during the transactional execution (step 802). This can involve either clearing or simply ignoring register file changes made during transactional execution. This is easy to accomplish because the old register values were checkpointed prior to commencing transactional execution. The system also clears load marks from cache lines in L1 data cache 115 (step 804), and drains store buffer entries generated during transactional execution without committing them to the memory hierarchy (step 806). At the same time, the system unmarks corresponding L2 cache lines. Finally, in one embodiment of the present invention, the system branches to a target location specified by the STE instruction (step 808). The code at this target location attempts to re-execute the critical section as is described above with reference to step 314 of FIG. 1.

STE Instruction

[0081] Referring to FIG. 9, STE instruction 900 marks the beginning of a block of instructions to be executed transactionally. When STE instruction 900 is encountered during program execution, the system commences transactional execution of the block of instruction immediately following STE instruction 900. Changes made during this transactional execution are not committed to the architectural state of the processor until the transactional execution successfully completes, at which point all of the changes are committed in one atomic operation as described above with reference to FIG. 7.

[0082] Note that STE instruction 900 includes an operation code (op code) 902, which identifies the STE instruction and a branch target 903. Branch target 903 contains a PC-relative branch target that specifies where to jump to when a failure occurs during transactional execution.

5 [0083] Instead of jumping to a branch target during a failure, the system can alternatively take other actions, such as automatically attempting the re-execute the block of instructions, or attempting to acquire a lock on the block of instructions and then re-executing the block of instructions. The system can also set state information within the processor to indicate a failure during transactional
10 execution of the block of instructions. This enables other software executed by the processor to manage the failure.

[0084] STE instruction 900 is somewhat unique in that its semantics are defined by what happens after the instruction, in a following transactionally executed block of code. If a commit instruction is successfully reached, STE
15 instruction 900 behaves as if it were a NOOP. In this case, the code in the critical section following STE instruction 900 is successfully completed, and execution resumes with the code following the commit instruction.

[0085] On the other hand, if a commit instruction is not successfully reached, the STE instruction behaves as if it were an annulling “branch always”
20 instruction that causes a status register to be updated with a value indicating the reason for failure. In this case, no visible effect remains from the instructions in the unsuccessful critical section.

[0086] Failures can arise from a number of sources. For example, in one embodiment of the present invention, transactional execution is considered to
25 have failed if any of the following events occur before reaching a commit instruction: (1) a fail instruction is reached; (2) an instruction not allowed in critical sections is encountered, such as FLUSH, MEMBAR, DONE, RETRY,

and any write of an ASR register, privileged register or internal ASI register; (3) a trap occurs; (4) an interrupt occurs; (5) a coherent load or store operation from another processor may have caused the sequence of loads and stores within the critical section to not appear atomic to other processors; (6) the processor runs out of hardware resources to buffer store operations from the critical section; (7) the processor runs out of hardware resources to track load operations from the critical section; and (8) too many stalling operations are encountered. This can be caused, for instance, by too many operations generating cache misses.

10 **Commit Instruction**

[0087] The commit instruction 910 marks the end of a block of instructions to be executed transactionally. Commit instruction 910 is used in conjunction with STE instruction 900 to delineate a block of instructions to be transactionally executed, such as a critical section. When commit instruction 910 is encountered during execution of a program, the system completes transactional execution of the block of instruction immediately preceding the commit instruction. This involves atomically committing changes made during the transactional execution to the architectural state of the processor as is described above with reference to FIG. 7. The system also resumes normal non-transactional execution.

[0088] If a commit instruction is reached, and there was no preceding STE instruction, the commit instruction is treated as a NOOP.

[0089] Referring to FIG. 9, the commit instruction 910 includes an op code 912, which identifies the instruction, and no operand is required.

25 [0090] In a variation of the commit instruction, after committing changes made during the transactional execution, the system resumes transactional execution of instructions following the commit instruction. Hence, instead of

terminating transactional execution, the system commits the changes made during transactional execution thus far and resumes transactional execution.

Fail Instruction

5 **[0091]** Fail instruction 920 terminates transactional execution without committing results of the transactional execution to the architectural state of the processor. Hence, fail instruction 920 directs the hardware to revert back to the program state at the point of the preceding STE instruction. This involves discarding changes made during transactional execution as is described above
10 with reference to FIG. 8.

[0092] After the changes are discarded, the system can perform a number of different actions. In one embodiment of the present invention, the system branches to a branch target specified by a corresponding start transactional execution (STE) instruction. In another embodiment, the system branches to a
15 location specified by the fail instruction itself, such as alternative branch target 923 in fail instruction 920 in FIG. 3. In yet another embodiment, the system attempts to re-execute the block of instructions.

[0093] Note that instead of immediately causing the transactional execution to fail, the fail instruction can simply set state information within the
20 processor to indicate that a failure has occurred, and the actual failure actions can take place later, for example upon encountering a subsequent commit instruction.

[0094] If a fail instruction is reached, and there was no preceding STE instruction, the fail instruction is treated as a NOOP.

[0095] Referring to FIG. 9, the fail instruction 920 includes an op code
25 912, which identifies the instruction, and possibly an alternative branch target 923.

[0096] Note that to facilitate “transactional execution” software can specify a speculative region of code, wherein selected memory accesses may be committed atomically to the architectural state of the processor as a transaction.

[0097] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.